

A Human-in/on-the-Loop Framework for Accessible Text Generation

Lourdes Moreno, Paloma Martínez

Universidad Carlos III de Madrid
Av. Universidad, 30, Leganés, 28911, Spain
lourdes.moreno@uc3m.es, paloma.martinez@uc3m.es

Abstract

Plain Language and Easy-to-Read formats in text simplification are essential for cognitive accessibility. Yet current automatic simplification and evaluation pipelines remain largely automated, metric-driven, and fail to reflect user comprehension or normative standards. This paper introduces a hybrid framework that explicitly integrates human participation into LLM-based accessible text generation. Human-in-the-Loop (HiTL) contributions guide adjustments during generation, while Human-on-the-Loop (HoTL) supervision ensures systematic post-generation review. Empirical evidence from user studies and annotated resources is operationalized into (i) checklists aligned with standards, (ii) Event-Condition-Action trigger rules for activating expert oversight, and (iii) accessibility Key Performance Indicators (KPIs). The framework shows how human-centered mechanisms can be encoded for evaluation and reused to provide structured feedback that improves model adaptation. By embedding the human role in both generation and supervision, it establishes a traceable, reproducible, and auditable process for creating and evaluating accessible texts. In doing so, it integrates explainability and ethical accountability as core design principles, contributing to more transparent and inclusive NLP systems.

Keywords: text accessibility, simplification, human-in-the-loop validation, LLM-based generation

1. Introduction

Access to understandable information is a fundamental right, particularly in critical domains such as health, education, legal, and public administration. Many people face barriers when interacting with complex or technical texts, including people with low literacy, reading comprehension difficulties, cognitive impairments, older adults, and non-native speakers (OECD, 2024). To address these challenges, two accessibility-oriented content formats have gained prominence: Plain Language (PL), codified in ISO 24495-1:2023 (ISO, 2023), and Easy-to-Read (ER), defined by standards such as the Spanish standard UNE 153101:2018 (UNE, 2018) and international guidelines (IFLA, 2010; Inclusion Europe, 2025).

Recent advances in Natural Language Processing (NLP), particularly with Large Language Models (LLMs), have enabled automatic text simplification at lexical, syntactic, and discourse levels (Vaswani et al., 2017; Nisioi et al., 2017; Shardlow et al., 2024). Nevertheless, evaluating and ensuring the accessibility of such outputs remains a major challenge. Widely used metrics such as BLEU (Papineni et al., 2002), SARI (Xu et al., 2016), SAMSA (Sulem et al., 2018), or BERTScore (Zhang et al., 2019) capture partial aspects of quality-lexical simplicity, structural fidelity, and semantic similarity-but fail to reflect user comprehension (Alva-Manchego et al., 2021).

Parallel to advances in NLP, the Human-Computer Interaction (HCI) community has emphasized participatory and empirical validation methods for accessibility (Lazar et al., 2017). Research

described in (Moreno et al., 2024; Alarcon et al., 2024, 2023) combines NLP resources with user-centered evaluations providing both methodological insights and empirical evidence on how cognitive accessibility can be implemented in practice.

However, most current approaches rely either on automatic metrics or ad hoc human validation, rarely integrating both into a systematic workflow. In the field of NLP and human-AI interaction, this integration has been conceptualized through Human-in-the-Loop (HiTL) and Human-on-the-Loop (HoTL) approaches. HiTL involves human input during the training or generation process, e.g., selecting or editing candidate simplifications, or guiding the model in fine-tuning (Amershi et al., 2014; Zanzotto, 2019). HoTL, by contrast, emphasizes post-hoc supervision, where humans oversee system outputs through structured review, guidelines, or monitoring mechanisms (Shneiderman, 2020). While HiTL has been applied to tasks such as interactive summarization or controllable text generation, and HoTL to safety monitoring in AI systems, these paradigms have not yet been fully exploited for cognitive accessibility in text simplification.

To fill the gap between automatic simplification and actual cognitive accessibility, we propose a framework that makes human participation a core design principle. HiTL enables interactive guidance during generation, such as accepting or rejecting lexical substitutions, while HoTL supervision provides structured oversight through checklists, trigger rules, and accessibility key performance indicators (KPIs) after outputs are produced. Together, these mechanisms ensure that accessible text generation is not only driven by metrics but also

grounded in human judgments, standards, and empirical validation.

At a glance, the framework combines: (i) HiTL guidance during generation, (ii) automatic metric-based checks, (iii) Event–Condition–Action (ECA) triggers for escalation, (iv) HoTL expert review, and (v) adaptation through structured feedback.

The remainder of this paper is organized as follows: Section 2 reviews related work on language accessibility standards, Section 3 reviews previous work on text simplification evaluation and HCI methods, and Section 4 presents background resources and empirical evidence. Section 5 introduces the proposed HiTL/HoTL framework (workflow, checklist, triggers, KPIs, integration). Section 6 discusses ethical and governance issues, and Section 7 concludes with a discussion and future work.

2. Plain Language (PL) and Easy-to-Read (ER) Standards

For NLP tools to be effective and socially relevant, they must be aligned with PL and ER standards, which promote the adaptation of texts to the real needs of their users, especially those with intellectual disabilities (ID) or reading limitations.

The main international reference is ISO 24495-1:2023, which defines general principles applicable to any language and domain, establishing that texts must be easy to find, understand, use, and evaluate by their target audience (ISO, 2023).

At the European level, the European Strategy on the Rights of People with Disabilities 2021–2030 highlights the importance of access to understandable information for persons with disabilities, recognizing ER as a key tool for achieving cognitive accessibility. Similarly, Directive (EU) 2016/2102 requires public sector websites and mobile applications to be accessible in accordance with standards such as the European standard EN 301 549 (ETSI, 2021) and the WCAG web content standard (W3C, 2025), which includes requirements on textual comprehensibility.

In Spain, the UNE 153101 EX:2018 standard on *Lectura Fácil* (LF) stands out, establishing guidelines for the drafting, layout, and validation of texts, emphasizing the participation of people with ID in the evaluation process (UNE, 2018). There are other national standards and legal frameworks that also shape PL/ER practice. Examples include ÖNORM V2100:2017 in Austria (administrative texts in LF), DIN SPEC 33429 in Germany (PL criteria for public documents), and the Equality Act (2010) in the United Kingdom, which requires public and private organizations to ensure accessibility, including clear information for people with cognitive disabilities. In addition, the UK General Data Pro-

tection Regulation (UK GDPR) and the Data Protection Act (2018) require privacy-related information to be clear and understandable. In healthcare, the NHS England Accessible Information Standard (AIS) requires accessible formats such as ER for people with cognitive or sensory disabilities. The Plain English Campaign certifies texts with the Crystal Mark, a symbol of linguistic clarity, while the GOV.UK Style Guide mandates the use of clear language on government digital platforms (PLAIN, 2021).

Although PL and ER share the overall goal of making text understandable, there are important differences. PL is aimed at a broad audience and seeks to reduce unnecessary technical terms and complex expressions. It typically relies on brevity, inclusive everyday language, and a structure that allows key information to be easily located, using direct style, active voice, and short paragraphs. ER is specifically designed for people with reading comprehension difficulties, such as people with ID, and therefore requires stricter guidelines for simplification and content presentation: bulleted lists, high-frequency words, abundant use of visual examples, and single-idea sentences. In addition, verification with real readers with ID is essential to validate ER texts (IFLA, 2010; Inclusion Europe, 2025; UNE, 2018).

The core principles shared by PL and ER include avoiding information overload, using familiar examples and metaphors, maintaining consistency in terminology, and reinforcing content with visual elements that facilitate comprehension. These guidelines constitute a valuable source of best practices that NLP developers can implement in automatic text simplification systems: controlling sentence length, adapting style to the target audience, and validating outputs with real users to ensure that simplifications effectively improve comprehension, usability, and cognitive accessibility.

3. Related Work

Research on text simplification and accessibility has evolved along three main lines: automatic evaluation metrics, HCI methods for cognitive accessibility, and human participation in NLP workflows. These strands converge in the challenge of ensuring that automatically simplified outputs are not only linguistically well-formed, but also cognitively accessible and aligned with normative standards.

Early work on automatic metrics focused on surface-level overlap with reference corpora, such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004). While useful in machine translation, these metrics have drawbacks as they penalize valid reformulations and ignore structural or semantic changes (Xu et al., 2016; Sulem et al., 2018). Ded-

icated metrics such as SARI (Xu et al., 2016) and DSARI (Maddela et al., 2021) improved lexical coverage by rewarding additions, deletions, and re-tentions, but still fail to capture global readability and meaning preservation. Structural metrics like SAMSA (Sulem et al., 2018) focus on sentence splitting and discourse structure, while semantic similarity methods such as BERTScore (Zhang et al., 2019) and MoverScore (Zhao et al., 2019) leverage embeddings to better capture meaning preservation. More recently, LENS (Maddela et al., 2022) has demonstrated a higher correlation with human judgments by learning evaluation from annotated data. However, systematic reviews conclude that no single metric fully captures the multifaceted requirements of accessibility (Alva-Manchego et al., 2021).

In parallel, the HCI community has developed methods to assess and design for cognitive accessibility. Standards (ISO, 2023; W3C WAI, 2021, 2023; ETSI, 2021) emphasize clarity, structure, and validation with end users. Empirical studies show that accessible interfaces can enhance comprehension when supported with contextual glossaries, pictograms, and multimodal cues. Alarcon et al. (2024); Moreno et al. (2024) developed an approach that combined corpus-based resources with participatory evaluations, validating glossary-based user interfaces with users with diverse cognitive profiles and developing annotated corpora for lexical simplification, (Alarcon et al., 2023). This line of work illustrates how user-centered design and linguistic resources can be integrated into accessibility evaluation. Works including user evaluations of simplification approaches (Ledoyen et al., 2025) proposed an evaluation framework in which human assessments supplement automatic metrics. Goldsack et al. (2023) evaluated a system that generates summaries of biomedical articles for lay readers, focusing on readability and factual-ity; the evaluation was conducted by NLP experts through questionnaires, leaving lay users aside. Other approaches simulate human behavior using multi-agent systems, such as ExpertEase (Mo and Hu, 2024), where multiple agents cooperate on text simplification tasks using external tools and guidelines, but no human evaluation is performed.

Finally, recent work has explored human participation in NLP architectures. HiTL approaches are widely used in annotation and training pipelines, allowing humans to refine data, correct model outputs, or steer generation in real time (Chiang and Lee, 2023). Gao et al. (2025) describe an evaluation of an LLM-based simplification system aligned using Direct Preference Optimization (DPO), using human feedback from people with ID (15 participants) and expert evaluators. The final evaluation indicated the effectiveness of DPO for personal-

izing LLM-based simplification, but also points to challenges such as improved HCI methods for eliciting preferences. More recently, Ilgen and Hattab (2025) defined a human-centered readability score based on a five-dimensional framework (tone, trust, cultural relevance, feasibility, and clarity), integrating automatic metrics with human experience around readability.

HoTL approaches emphasize structured supervision and post-hoc control, in which humans monitor, validate, or adjust system outputs through review procedures, guidelines, or oversight mechanisms (Li et al., 2020; Shneiderman, 2020; Carrer et al., 2024). In the context of text simplification, these paradigms remain underexplored, although related work in summarization and factual consistency shows the potential of combining automatic scoring with human oversight (Manakul et al., 2023). This convergence reinforces the shift toward hybrid workflows that systematically integrate human validation within NLP evaluation pipelines.

Recent work has also addressed the trustworthiness and safety of text simplification. Hayakawa et al. (2025) proposed an efficient framework for lexical simplification with small LLMs, emphasizing privacy, efficiency, and the detection of harmful simplifications through probabilistic safety signals. Their approach highlights the need for traceable and accountable simplification pipelines.

In summary, existing metrics, HCI methods, and human-in/on-the-loop frameworks each addressing parts of the problem but rarely intersecting. What is missing is a unified approach that reuses language resources, incorporates HCI evidence, and defines hybrid workflows where humans complement automation in the generation, evaluation, and adjustment of PL and ER outputs. This gap motivates the approach proposed in this paper.

4. Background Resources

Evidence is drawn from user studies and corpora developed in previous research on cognitive accessibility. Data from older adults, people with intellectual disabilities (ID), and expert annotators are used in tasks such as complex word identification (CWI), synonym acceptance, corpus annotation, and User Interface (UI) validation. Only selected findings are summarized here, providing preliminary evidence to motivate a hybrid HiTL-HoTL framework and serving as the basis for accessibility KPIs, review triggers, and human intervention points.

4.1. Complex Word Identification Task

Alarcon et al. (2024) did an evaluation for CWI task where 50 participants (25 older adults and 25 with ID) were asked to highlight words considered diffi-

cult in domain-specific texts. In this study, participants directly marked the words they perceived as difficult, and system behavior was analyzed through standard CWI metrics as user-grounded evidence of perceived lexical difficulty.

Results indicated systematic differences between groups. Older adults tended to select a larger number of candidates, which led to higher recall (0.73) but lower precision (0.54). By contrast, participants with ID marked fewer items, resulting in higher precision (0.58) and a lower F1 score (0.56).

Within each group, educational background also influenced performance: older adults with higher education showed a better balance between recall and precision, whereas participants with lower literacy showed greater variability in their selections. These complementary profiles suggest that while older adults provide broader coverage, participants with ID contribute more selective judgments. This divergence motivates adaptive HiTL mechanisms that can balance broader lexical coverage against more selective judgments across user profiles.

4.2. Synonym Acceptance Task

Alarcon et al. (2024) carried out an experiment with the same 50 participants (25 older adults and 25 with ID), who were asked to judge the adequacy of proposed synonyms to complex words in the context. Results indicate profile-dependent differences across groups (Table 1). Older adults tended to reject all candidates more frequently (16.6%) than participants with ID (7.6%). In contrast, participants with ID showed a stronger tendency to accept at least one synonym (67.5% vs 58.9%), while full acceptance was very similar across groups (24.5% vs 24.9%). These patterns suggest that participants with ID may be more flexible in their judgments, whereas older adults are more conservative.

Beyond group-level differences, educational background also influences synonym acceptance. Participants with lower education levels (primary or no studies) achieved higher rates of correct acceptance ($\approx 78\text{--}80\%$), while those with higher education were more likely to partially accept or reject candidates (e.g., $\approx 65\%$ correct with secondary education). This effect was statistically significant, confirming that literacy levels shape how simplifications are perceived and judged.

Taken together, these results indicate that synonym acceptance is not homogeneous across user profiles, and that both group characteristics (older vs. ID) and individual factors (such as education level) must be considered when defining acceptance thresholds. For HiTL workflows, this evidence motivates adaptive mechanisms sensitive to user profile; for HoTL supervision, it suggests the need to monitor systematic rejection patterns as triggers for expert review.

Group	None	Some	All
Older adults	16.6%	58.9%	24.5%
Participants with ID	7.6%	67.5%	24.9%

Table 1: Synonym acceptance rates by participant group ($n = 50$). *None/Some/All* = none accepted / some accepted / all accepted.

4.3. Expert Annotation Reliability

Concerning resources, EASIER corpus, (Alarcon et al., 2023), was annotated by three expert linguists and is composed of 260 documents with 8,155 words labelled as complex words and 5,130 words with at least one proposed context-aware synonym associated. This corpus provides evidence about the reliability of lexical complexity judgments. Agreement was higher for multi-word expressions ($\kappa = 0.64$) than for single words ($\kappa = 0.52$), suggesting that multi-words may be more stable units for supervised training and for HoTL supervision. These results provide complementary evidence to user studies and support the integration of supervised resources into hybrid evaluation workflows.

4.4. HCI-based Validation Methods

In addition to corpus annotation, (Moreno et al., 2024) describes a system where glossaries of definitions, synonyms, and pictograms are integrated into cognitively accessible user interfaces, which were validated with end users, including older adults and people with ID. Validation relied on HCI-based methods applied directly with participants (Lazar et al., 2017), such as: (i) comprehension questionnaires measuring retention of key information; (ii) navigation tasks evaluating time and success rates to locate content in documents; (iii) glossary activation logs recording actual use of contextual support; and (iv) usability feedback collected through standardized instruments. These methods provided empirical evidence of how accessibility is perceived and implemented by real users, beyond what automatic metrics can analyze.

4.5. Resources for Language Model Adaptation

Beyond user studies and user interface validation, evidence also comes from resources applied in simplification system development in evaluation campaigns. Participation in a shared task on text simplification (Moreno et al., 2025) demonstrated that widely used automatic simplification metrics (e.g., BLEU, SARI, FKGL) fail to capture comprehension and accessibility principles, requiring complementary human feedback for calibration. In parallel, (Martínez et al., 2024) emphasized the need to

align evaluation with international standards such as ISO 24495-1 and with empirical user evidence.

These works highlight two key points. First, resources such as annotated corpora, synonym judgments, and comprehension tests are not only valuable for benchmarking but can also inform prompt design and fine-tuning of LLMs. For instance, simplification rules derived from corpus annotation can be encoded as constraints in prompting strategies. Second, iterative metrics-based evaluation cycles followed by human review and adjustment constitute an implicit HiTL–HoTL workflow, showing how human participation can guide and supervise LLM adaptation.

These resources also anticipate integration into adaptive training pipelines in the future. Preference signals such as synonym acceptance thresholds, comprehension test scores, and glossary activation logs can be repurposed not only to evaluate outputs but also to supervise iterative fine-tuning cycles. This perspective directly links empirical evidence with emerging preference-based approaches to LLM alignment. In this sense, accessibility-grounded corpora provide domain-specific signals that can complement large-scale, generic datasets in both prompting strategies and preference-based adaptation.

4.6. Implications for HiTL–HoTL Workflows

Taken together, these results illustrate how human participation most effectively complements automation. Older adults provide coverage in CWI, people with ID contribute precision, expert annotation stabilizes supervision via multi-word units, and validated interface patterns ensure delivery through accessible modalities. Moreover, resources for model adaptation (Martínez et al., 2024) and evidence of metric limitations (Alva-Manchego et al., 2021; Moreno et al., 2025) indicate that human input must not only validate but also guide training and prompting strategies.

This convergence of corpus-based evidence, user studies, and model adaptation resources sets the stage for a hybrid framework. Section 5 implements these insights by defining a Human-in/on-the-Loop (HiTL–HoTL) workflow structured around checklists, trigger rules, and KPIs.

5. Proposed Framework

We propose a hybrid framework that integrates Human-in-the-Loop (HiTL) and Human-on-the-Loop (HoTL) mechanisms across accessible text generation, metric-based evaluation, and adaptation of PL/ER outputs. Grounded in the evidence summarized in Section 4 (CWI, synonym accep-

tance, expert annotation, UI validation, and model adaptation), the framework operationalizes this evidence through checklists, ECA trigger rules, and profile-specific KPIs that govern the workflow depicted in Figure 1. The contribution is therefore an operational and traceable framework that consolidates prior empirical evidence, annotated resources, and accessibility-oriented supervision into a unified workflow.

5.1. Overall Workflow

At a high level (Figure 1), the workflow is a multi-stage, iterative pipeline. Evidence and resources (A) condition generation with HiTL (B) through inputs and constraints; the output plus interaction logs are then automatically evaluated (C). From (C), metrics feed rule-based checking (D) and quality indicators (E) that determine routing: if rule conditions are met, the ECA rule engine activates HoTL supervision (F); if quality trends or thresholds suggest systematic issues, signals are sent to Adaptation and Learning (G). Expert editing operations from (F) are consolidated in (G) and returned to (B) as updates to prompts/policies or the model, closing the loop; once an approved decision is issued in (F), the result is delivered via an accessible User Interface (H).

Evidence from Section 4 indicates that differences in CWI between older adults and people with ID inform HiTL checkpoints; profile-specific synonym-acceptance patterns motivate HoTL escalation; higher agreement in multi-word annotations defines stable units for supervision; and user-validated resources can be embedded into prompts and fine-tuned models for adaptation.

The following sections provide the technical details: Section 5.2 on HiTL adjustments, Section 5.3 on the HoTL checklist, Section 5.4 on rule-based routing, Section 5.5 on quality indicators, and Section 5.6 on how these signals are leveraged for adaptation.

5.2. Human-in-the-Loop (HiTL) Adjustments

HiTL participation is enacted during generation (Figure 1, B). Evidence from Section 4 indicates broader coverage (higher recall) for older adults and higher precision for participants with ID; these complementary profiles are implemented through interactive steps that (i) request short, active sentences, (ii) discourage jargon and expand acronyms, (iii) split long sentences, and (iv) fix meaning distortions while preserving facts and numeric integrity. User-validated resources (Section 4), annotation guidelines, synonym judgments, and glossary constraints are encoded directly into

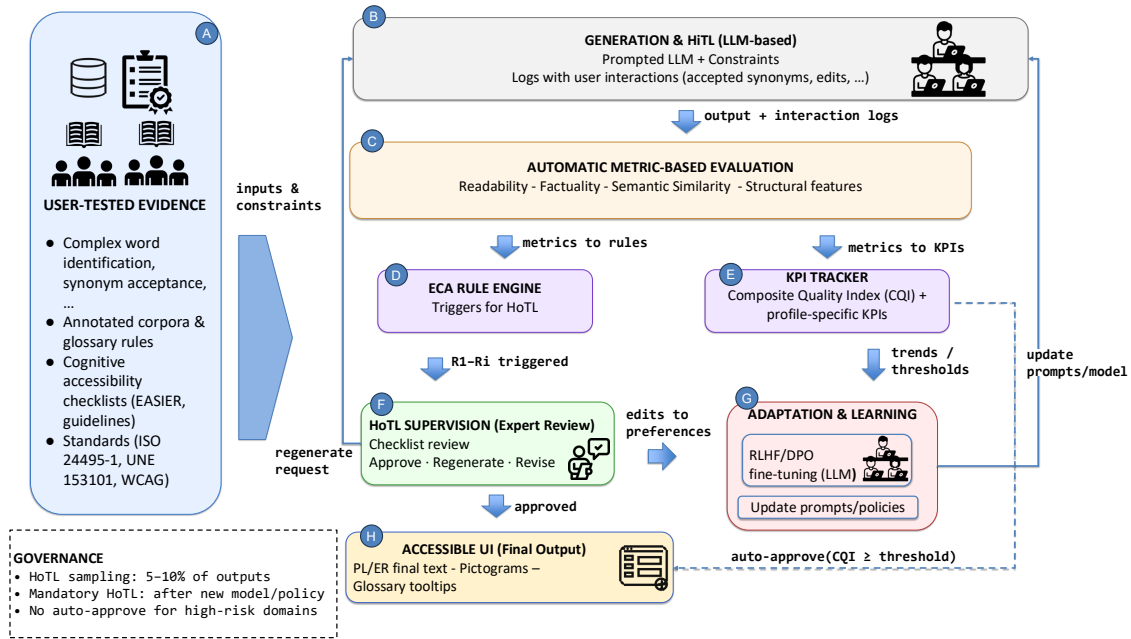


Figure 1: Overall workflow of the proposed HiTL-HoTL framework.

prompts or, when needed, into fine-tuning procedures. Interaction traces, produced at generation time, feed automatic evaluation (Figure 1, B→C), while stabilized guidance (e.g., glossary rules) is later consolidated in adaptation and learning and pushed back to generation (Figure 1, G→B).

5.3. Human-on-the-Loop (HoTL) Checklist

While HiTL constrains and guides generation under uncertainty, HoTL provides accountable human adjudication when borderline, risky, or conflicting cases cannot be resolved with sufficient confidence through automatic checks.

HoTL supervision provides a structured review layer when automatic metrics do not guarantee comprehension. The checklist integrates international standards (ISO, 2023; UNE, 2018) with evidence from synonym acceptance, expert-annotation reliability, interface validation, and model-adaptation studies (Section 4).

Table 2 summarizes how accessibility dimensions map to concrete checklist items with human-centered validation (e.g., comprehension tests, reading times, usability questionnaires), including a row on prompt/model adaptation where user-validated resources (corpus rules, glossary constraints, synonym judgments) can be encoded into prompt design or fine-tuning.

Outputs are considered checklist-compliant when at least two-thirds of the dimensions are satisfied. In the workflow, expert review is reached when

rule checks route an item to HoTL (Figure 1, D→F); edits inform adaptation for subsequent cycles (Figure 1, F→G), regeneration can be requested where appropriate (Figure 1, F→B), and delivery follows an approved decision (Figure 1, F→H).

5.4. Trigger Rules for Review

Systematic synonym-rejection patterns (Section 4) and divergences between readability and comprehension (Section 4) motivate ECA rules that determine when HoTL supervision is required. Following automatic evaluation (Figure 1, C), metric outputs feed the rule engine; when trigger conditions are met, cases are escalated to expert review (Figure 1, D→F). Listing 1 provides illustrative examples of accessibility-related risks that may trigger escalation, with thresholds that can be made profile-specific.

Listing 1: ECA rules for accessibility-driven triggers (conditions reference metric thresholds).

```

RULE R1 // fluency over meaning
IF Readability_FH > 80 AND BERTScore < 0.85
THEN Activate HoTL supervision

RULE R2 // excessive deletion of essentials
IF SARI_deletions > 0.40 AND AlignScore < 0.80
THEN Activate HoTL supervision

RULE R3 // structural clarity compromised
IF DSARI < theta_DSARI OR SAMSA < theta_SAMSA
THEN Activate HoTL supervision

```

Thresholds (e.g., 80 for readability, 0.85 for BERTScore) are calibrated per profile and domain, and can be revised as evidence is gathered. Addi-

Dimension	Operational check	Human-centered validation
Lexical clarity	Common words; explain acronyms	Comprehension tests
Syntactic simplicity	≤20 words per sentence; one idea per sentence	Eye-tracking; reading times
Structural clarity	Logical order; headings; lists	Task success rates
Relevance	Preserve essentials; avoid redundancy	User judgments
Multimodal support	Glossaries; pictograms	Usability questionnaires; user testing
Prompt/model adaptation	Encode rules, glossary constraints, and domain specificity	HiTL prompt design; fine-tuning

Table 2: HoTL checklist integrating accessibility dimensions with human-centered validation.

tional rules may include violations of terminology/glossary constraints or missing mandatory sections, ensuring that escalation to HoTL remains sensitive to accessibility-critical failures.

5.5. Key Performance Indicators (KPIs)

KPIs are defined to connect evaluation with the user evidence in Section 4. Examples include comprehension gains from user tests, synonym-acceptance rates stratified by profile, glossary-activation rates in the UI, the recall-precision balance in CWI, and improvements in model-adaptation accuracy.

To implement these indicators, ECA rules are specified in Listing 2. Listing 2 summarizes the KPI layer used for routing and monitoring. In addition, a Composite Quality Index (CQI) integrates Readability (0.4), Semantic Fidelity (0.3), and Structural Clarity (0.3), with weights summing to 1 and intended as an initial operational weighting for routing decisions. In the workflow (Figure 1), KPIs populate (E) and support monitoring, thresholding, and adaptation decisions together with rule checks.

Listing 2: Accessibility-oriented KPIs and composite quality index.

```

KPI_1: IF comprehension_test_score(user) >=
  baseline + delta
  THEN record comprehension_gain

KPI_2: IF synonym_acceptance_rate(profile) >=
  theta_profile
  THEN validate lexical_clarity

KPI_3: IF glossary_activation_rate >= tau
  THEN confirm multimodal_support_used

KPI_4: IF recall_precision_balance(CWI) within
  [alpha, beta]
  THEN accept CWI_performance

KPI_5: IF model_adaptation_accuracy >=
  previous_cycle + epsilon
  THEN record adaptation_improvement

Composite_Quality_Index (CQI):
  CQI = 0.4*Readability + 0.3*Semantic_Fidelity
  + 0.3*Structural_Clarity

Decision threshold: CQI >= gamma (calibrated to
  comprehension outcomes)

```

5.6. Integration of Evidence

Beyond evaluation, the proposed KPIs and triggers can be implemented as structured preference signals for adaptation. While approaches such as Reinforcement Learning from Human Feedback (RLHF) (Stiennon et al., 2020), Direct Preference Optimization (DPO) (Rafailov et al., 2023), and other preference-based fine-tuning strategies (Ziegler et al., 2019; Böhm et al., 2019), including multi-agent feedback frameworks for adaptive simplification (Mo and Hu, 2024), do not strictly require domain-specific signals, accessibility-grounded KPIs provide a practical basis to ensure that adaptation is aligned with PL and ER principles rather than generic preferences. In this view, synonym-acceptance judgments, comprehension outcomes, and UI usability/usage signals extend beyond measurement and drive iterative improvement (Figure 1, F→G→B).

To make this actionable, what to measure (KPIs) and what to do (triggers) are encoded in a unified ECA policy that governs routing and adaptation; Listing 3 abstracts both into a unified ECA policy.

Listing 3: Unified ECA formalization of KPIs and triggers.

```

# KPI definition: measurable function of metrics
  and/or user signals
KPI_i = f(metrics, user_signals)

# Trigger: ECA rule referencing one or more KPIs
TRIGGER_j:
IF KPI_i < theta_i OR combine(KPI_m, KPI_n) <
  threshold
THEN Activate HoTL supervision

```

5.7. Illustrative Framework Execution

To illustrate how the framework operates in practice, Table 3 presents a compact lexical simplification case derived from the user-study materials discussed in Section 4. The example shows how a potentially difficult term may motivate simplification, while still requiring escalation when a candidate substitutes risks changing the register or reducing domain adequacy.

This case illustrates a typical accessibility-sensitive situation for the proposed framework. The

Stage	Excerpt / signal	Framework interpretation
Original excerpt	“La caza y la pesca en la Comunidad de Madrid están sujetas a regulación especial y requieren una serie de trámites , entre ellos la obtención de licencia ...” / <i>“Hunting and fishing in the Community of Madrid are subject to special regulation and require a series of administrative procedures, including obtaining a licence ...”</i>	The source text contains a potentially difficult domain-specific term in an administrative context.
Initial adapted output	“...requieren una serie de papeleos ...” / <i>“...require a series of paperwork ...”</i>	The output simplifies the term, but may introduce an overly colloquial register and reduce administrative precision.
Detected issue/trigger	Readability may improve at the lexical level, but adequacy risk remains for domain register and terminology. / <i>Possible mismatch between lexical simplicity and administrative appropriateness.</i>	Trigger rule routes the case to HoTL review because lexical simplification alone does not guarantee appropriate wording.
HoTL-guided revision	Review the candidate substitute against the checklist dimensions for terminology, clarity, and domain adequacy. / <i>Check whether the simplified term remains precise, understandable, and appropriate for the target context.</i>	Checklist dimensions affected: lexical clarity, terminology consistency, and appropriateness of register.
Approved output	“...requieren varios trámites administrativos , entre ellos obtener una licencia ...” / <i>“...require several administrative procedures, including obtaining a license ...”</i>	Meaning-preserving simplification approved; the wording is clearer while maintaining administrative precision.

Table 3: Illustrative execution of the proposed framework on a lexical simplification case. Spanish excerpts are followed by an English gloss in italics.

original term (*trámites*) may motivate lexical simplification for some readers, but not every candidate substitute is equally appropriate. A form such as *papeleos* may appear simpler, yet it introduces a more colloquial register and weakens the administrative precision and pragmatic appropriateness of the original wording in this document type. In such cases, automatic simplification may improve local readability while still requiring HoTL review to verify domain adequacy, terminology consistency, pragmatic fit, and appropriateness for the intended user profile. The revised version therefore preserves accessible wording without sacrificing the intended administrative meaning.

5.8. Multi-agent Workflow Development

As a practical implementation pathway, the framework could be developed as a multi-agent workflow, with specialized components for evidence retrieval, generation, evaluation, rule checking, KPI tracking, and human review. In this view, an orchestrator coordinates the interaction between generation, automatic assessment, escalation rules, and HoTL supervision, while preserving the traceability of intermediate decisions. Each agent performs a distinct role within the human-in/on-the-loop workflow, supporting traceability, explainability, and adaptive control. Taken together, these agents form a traceable and adaptive workflow in which automation and human supervision jointly support accessible,

explainable, and accountable text generation. This architecture-specific realization is secondary to the present contribution, which remains the framework itself.

6. Ethical Considerations and Governance

The framework is designed for accountability. Every decision point is traceable via versioned prompts and policies, rule evaluations, KPI snapshots, and HiTL/HoTL logs, enabling post-hoc audit and model-update provenance. Explainability is supported by explicit ECA rules (why a route was chosen), checklist rationales (why edits were made), and profile-specific thresholds (why an item was approved or escalated). To protect users, logs follow data minimization and role-based access controls.

Governance safeguards ensure human oversight across the loop: HoTL sampling (5–10% of outputs) applies even under auto-approve; it is mandatory after new model or policy releases; and auto-approve is disabled for high-risk domains (e.g., dosages in drug treatments appearing in a discharge summary report, legal warnings in contracts). These controls and profile-specific thresholds uphold oversight while enabling continuous, evidence-based adaptation.

This governance setup is consistent with the

Ethics Guidelines for Trustworthy AI (European Commission, 2019), emphasizing human agency and oversight (HiTL/HoTL), transparency and explainability (machine-actionable rules with interpretable rationales), privacy and data governance (controlled retention and access), and fairness and accessibility (PL/ER principles). In addition, the proposed framework explicitly supports traceability and explainability as transversal safeguards: every generation, evaluation, and supervision step is logged and versioned, enabling end-to-end traceability from input prompts and human interventions to final outputs and model updates. Likewise, explainability is ensured by the declarative ECA rules and checklist rationales that make the decision path verifiable and intelligible both for experts and regulators. Together, these mechanisms reinforce accountability and foster confidence in the deployment of adaptive, human-supervised NLP systems for cognitive accessibility.

7. Conclusions

This paper proposes a hybrid framework that integrates Human-in-the-Loop (HiTL) and Human-on-the-Loop (HoTL) mechanisms to support cognitive accessibility in Plain Language and Easy-to-Read text generation. Grounded in empirical evidence from the EASIER project, expert annotation, and user interface validation, the framework translates accessibility principles into operational components such as checklists, trigger rules, and Key Performance Indicators (KPIs). These resources systematically complement automatic metrics with user-grounded evidence and expert supervision.

The framework builds on prior evidence suggesting complementary strengths across user groups, broader coverage from older adults and higher precision from people with intellectual disabilities, as well as the stabilizing role of expert annotations and the added value of interface validation for real-world deployment. Moreover, integrating accessibility-based corpora and evaluation signals illustrates how human feedback can guide prompt design, fine-tuning, and preference-based alignment strategies.

By linking accessibility KPIs and triggers with adaptation methods such as RLHF and DPO, the framework connects evaluation and training into a traceable, explainable, and reproducible workflow. These mechanisms strengthen accountability and transparency, ensuring that accessibility standards are embedded not only in model outputs but also in the generation and supervision processes.

Future work will extend validation across domains, explore quantitative measures of comprehension gain, and develop the multi-agent orchestration framework, expanding the roles of evidence,

evaluation, and reviewer agents to automate monitoring, explainability, and adaptive retraining. This direction aims to contribute to the standardization of accessibility-aware evaluation pipelines and to the creation of inclusive, transparent, and trustworthy NLP systems.

8. Acknowledgements

This work has been supported by grant PID2023-1485770B-C21 (Human-Centered AI: User-Driven Adapted Language Models-HUMAN_AI) by MICIU/AEI/10.13039/501100011033 and by FEDER/UE.

9. Bibliographical References

- Rodrigo Alarcon, Lourdes Moreno, and Paloma Martínez. 2023. [Easier corpus: A lexical simplification resource for people with cognitive impairments](#). *PLOS ONE*, 18(4):e0283622.
- Rodrigo Alarcon, Lourdes Moreno, Paloma Martínez, and José A. Macías. 2024. [Easier system. evaluating a spanish lexical simplification proposal with people with cognitive impairments](#). *International Journal of Human-Computer Interaction*, 40(5):1195–1209.
- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. [The \(un\)suitability of automatic evaluation metrics for text simplification](#). *Computational Linguistics*, 47(4):861–889.
- Saleema Amershi, Maya Cakmak, W. Bradley Knox, and Todd Kulesza. 2014. [Power to the people: The role of humans in interactive machine learning](#). *AI Magazine*, 35(4):105–120.
- Florian Böhm, Yang Gao, Christian M. Meyer, Ori Shapira, Ido Dagan, and Iryna Gurevych. 2019. [Better rewards yield better summaries: Learning to summarise without references](#). arXiv preprint arXiv:1909.01214.
- Luisa Carrer, Tatiana Anikina, Mounica Maddela, Fernando Alva-Manchego, and Carolina Scarton. 2024. [Towards holistic human evaluation of automatic text simplification](#). In *Proceedings of the 1st Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 66–79.
- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) arXiv preprint arXiv:2305.01937.

- ETSI. 2021. En 301 549: Accessibility requirements for ict products and services. https://www.etsi.org/deliver/etsi_en/301500_301599/301549.
- European Commission. 2019. Ethics guidelines for trustworthy ai. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.
- Y. Gao, K. Johnson, D. Froehlich, L. Carrer, and S. Ebling. 2025. Evaluating the effectiveness of direct preference optimization for personalizing german automatic text simplifications for persons with intellectual disabilities. arXiv preprint arXiv:2507.01479.
- T. Goldsack, Z. Zhang, C. Tang, C. Scarton, and C. Lin. 2023. Enhancing biomedical lay summarization with external knowledge graphs. arXiv preprint arXiv:2310.15702.
- A. Hayakawa, S. Bott, and H. Saggion. 2025. Towards trustworthy lexical simplification: Exploring safety and efficiency with small llms. arXiv preprint arXiv:2509.25086.
- HULAT Group. 2026. Artefacts for Irec 2026: A human-in/on-the-loop framework for accessible text generation. Zenodo record.
- IFLA. 2010. Directives for easy-to-read materials. <https://www.ifla.org/publications/directory-of-easy-to-read-publications/>
- B. İlgen and G. Hattab. 2025. Toward human-centered readability evaluation. arXiv preprint arXiv:2510.10801.
- Inclusion Europe. 2025. European standards and guidelines for easy-to-read information. <https://www.inclusion-europe.eu/easy-to-read-standards-guidelines/>.
- ISO. 2023. Iso 24495-1:2023 plain language—part 1: Governing principles and guidelines. <https://www.iso.org/standard/78907.html>.
- Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research Methods in Human-Computer Interaction*. Morgan Kaufmann.
- F. Ledoyen, G. Dias, A. Lechervy, J. Pantin, F. Maurel, Y. Chahir, E. Gouzonnat, M. Berthelot, S. Moravac, A. Altinier, et al. 2025. Inclusive easy-to-read generation for individuals with cognitive impairments. arXiv preprint arXiv:2510.00691.
- Nianyu Li, Adebola Adeyiga, and Rami Bahsoon. 2020. Explanations for human-on-the-loop: A probabilistic model checking approach. In *Proceedings of the IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, pages 181–187.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- M. Maddela, Y. Dou, D. Heineman, and W. Xu. 2022. Lens: A learnable evaluation metric for text simplification. arXiv preprint arXiv:2212.09739.
- Mounica Maddela, Fernando Alva-Manchego, and Wei Xu. 2021. Controllable text simplification with explicit paraphrasing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3536–3553, Online. Association for Computational Linguistics.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.
- Paloma Martínez, Alberto Ramos, and Lourdes Moreno. 2024. Exploring large language models to generate easy to read content. *Frontiers in Computer Science*, 6:1394705.
- K. Mo and R. Hu. 2024. Expertease: A multi-agent framework for grade-specific document simplification with large language models. *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9080–9099.
- Lourdes Moreno, Helen Petrie, Paloma Martínez, and Rodrigo Alarcon. 2024. Designing user interfaces for content simplification aimed at people with cognitive impairments. *Universal Access in the Information Society*, 23:99–117.
- Lourdes Moreno, Jesus M. Sanchez-Gomez, Marco Antonio Sanchez-Escudero, and Paloma Martínez. 2025. Hulat-uc3m @ clears2025 sub-task 1: Prompt-based simplification for plain language using spanish language models. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2025)*, volume 4098 of *CEUR Workshop Proceedings*. CEUR-WS.org, ISSN 1613-0073.
- Sergiu Nisioi, Sanja Štajner, Simone Paolo Ponzetto, and Liviu P. Dinu. 2017. Exploring neural text simplification models. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 85–91.

- OECD. 2024. *Skills Outlook 2024: Skills for Inclusive Transitions*. OECD Publishing.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- PLAIN. 2021. Plainlanguage.gov. <https://www.plainlanguage.gov>.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). arXiv preprint arXiv:2305.18290.
- Matthew Shardlow, Fernando Alva-Manchego, Riza Theresa Batista-Navarro, Stefan Bott, et al. 2024. [An extensible massively multilingual lexical simplification pipeline dataset using the multils framework](#). In *Proceedings of the 3rd Workshop on Tools and Resources for People with READING Difficulties (READI) @ LREC-COLING 2024*, pages 38–46.
- Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, 36(6):495–504.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2018. [Semantic structural evaluation for text simplification](#). arXiv preprint arXiv:1810.05022.
- UNE. 2018. Une 153101:2018 ex. lectura fácil. pautas y recomendaciones para la elaboración de documentos. <https://www.une.org/encuentra-tu-norma/busca-tu-norma/norma?c=N0060036>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.
- W3C. 2025. Wcag 2 overview. <https://www.w3.org/WAI/standards-guidelines/wcag/>.
- W3C WAI. 2021. Making content usable for people with cognitive and learning disabilities (coga). <https://www.w3.org/TR/coga-usable/>.
- W3C WAI. 2023. Web content accessibility guidelines (wcag) 2.2. <https://www.w3.org/TR/WCAG22/>.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Fabio Massimo Zanzotto. 2019. Human-in-the-loop artificial intelligence. *Journal of Artificial Intelligence Research*, 64:243–252.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [BERTScore: Evaluating text generation with BERT](#). arXiv preprint arXiv:1904.09675.
- W. Zhao, M. Peyrard, F. Liu, Y. Gao, C. M. Meyer, and S. Eger. 2019. [Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance](#). arXiv preprint arXiv:1909.02622.
- Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. [Fine-tuning language models from human preferences](#). arXiv preprint arXiv:1909.08593.

10. Language Resource References

A first set of reusable artefacts associated with the framework (checklist schema, ECA trigger templates, KPI definitions) is available through Zenodo ([HULAT Group, 2026](#)).

A. An Illustrative Lexical Simplification Example

This appendix provides the full passage underlying the illustrative example discussed in Table 3. The example is included to clarify how automatic checks, escalation rules, and HoTL review interact in the proposed workflow. It is not presented as a new evaluation result, but as a compact demonstration of framework execution on a lexical simplification case derived from the materials used in the user studies described in Section 4.

A.1. Original passage

“La caza y la pesca en la Comunidad de Madrid están sujetas a regulación especial y requieren una serie de **trámites**, entre ellos la obtención de licencia, para poder practicarlas.” / “*Hunting and fishing in the Community of Madrid are subject to special*

*regulation and require a series of **administrative procedures**, including obtaining a licence, in order to be practiced.”*

A.2. Initial adapted version

“La caza y la pesca en la Comunidad de Madrid tienen normas especiales y requieren una serie de **papeleos** para poder practicarlas.” / *“Hunting and fishing in the Community of Madrid have special rules and require a series of **paperwork** in order to be practiced.”*

A.3. Illustrative revision after HoTL review

“La caza y la pesca en la Comunidad de Madrid tienen normas especiales y requieren varios **trámites administrativos**, entre ellos obtener una licencia, para poder practicarlas.” / *“Hunting and fishing in the Community of Madrid have special rules and require several **administrative procedures**, including obtaining a licence, in order to be practiced.”*